

[System and Method to distribute reasoning and pattern matching in forward and backward chaining rule engines]

Abstract

Distributed pattern matching defines extensions and improvements to the RETE algorithm and to the forward chaining method of inference in rule systems. The same approach is envisaged for a goal-driven, backward chaining inference method where data is used to satisfy new goals in an interactive session between different engines, different systems, or between a system and a user interface. Utilizing techniques from Remote Procedure Calling (RPC), distributed processing, parallel processing and grid computing, the invention defines a novel algorithm to implement a system capable of distributed pattern matching. Reasoning over large data sets exceeding millions of rows in a reactive mode with a single engine is difficult due to hardware limitations. Using a cluster of inexpensive systems, reasoning over large datasets is more efficient and provides better scalability and response time. Distributed pattern matching is a significant innovation and enables rule engines to reason over large sets of data. Current rule

engines scale vertically by upgrading the hardware. Distributed pattern matching has no preference towards vertical or horizontal scaling. There are classes of problems that require a rule engine to reason over the entire database. This type of problem is difficult to solve with collaborative reasoning, where multiple symmetric systems are used to process discrete blocks of data. Distributed pattern matching overcomes the limitations of collaborative reasoning and helps solve a new class of problems. This innovation will also facilitate system integration, lower costs of maintenance in large organizations, and improve the performance and quality of applications in many areas, including: electronic procurement, e-government, real-time data management systems, insurance systems, enterprise resource management, and personalization, to just cite a few applications.